

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Automated Generator of Input-Validation Filters

Inventor(s):

Daniel Simon

Valentin Razmov

Jonathan Pincus

ATTORNEY'S DOCKET NO. MS1-744US

TECHNICAL FIELD

This invention generally relates to a technology for facilitating the automated generation of input-validation software filters.

BACKGROUND

In some instances, it can be beneficial to constrain an application's input to only that which meets specific criteria. Such input may be created directly by a person (i.e., a human) or may be the direct product of a computing component.

Input from Humans

Ordinary forms are designed for gathering specific input from a human. For example, if a form field requires a date, the only appropriate input is a valid date (such as June 20, 2001). Electronic forms exist (e.g., JetForm™), designed to limit the format and the type of input that a human enters into such forms. However, many such electronic forms are not well suited for collecting input from computing components.

Input from Computing Components

Many computing components accept input directly from other computing components. Even if the data provided by a computing component originates from a human, such data is considered input from a computing component when it is accepted directly from that component and not a human.

Broadly speaking, a computing component is one that excludes a direct human intervention. Examples of such computing components include (by way of

example, but not limitation): applications, program modules, Web pages, Web servers, client computers, and mobile devices.

A prime example of input from a computing component is input accepted by a Web server from another computer (e.g., a client) over a communications network (such as the Internet). Although the bulk of such input is benign and has valid formatting, sometimes such input is accidentally or purposefully made ill-formed. Such purposeful actions may be part of a malicious attack directed by one or more attackers (i.e., interlopers).

Malicious Attacks

Malicious attackers routinely gain unauthorized access into computing systems (e.g., Web servers) by exploiting poor or non-existent filtering of input from computing components.

A malicious interloper may attack a computer system by supplying input, which can manipulate the system into performing otherwise unauthorized actions targeted at subverting the integrity of the system. Examples of such malicious attacks include buffer overruns and other forms of invalid input.

More information on these types of attacks may be found in the following articles from CERT®:

- “Understanding Malicious Content Mitigation for Web Developers”, CERT Coordination Center, Feb. 2000, http://www.cert.org/tech_tips/malicious_code_mitigation.html and http://www.cert.org/tech_tips/malicious_code_FAQ.html; and

- “Malicious HTML Tags Embedded in Client Web Requests”, CERT Coordination Center, Feb. 2000, <http://www.cert.org/advisories/CA-2000-02.html>).

Buffer Overruns

To cause a buffer overrun, an attacker sends much more data to an application (such as a Web application) than the application can handle in its buffer. Here, “much more” means more data in size than the application implicitly assumes it should be receiving. If the application does not check the length of the data received, it may simply copy it into a fixed-size array of bytes on the stack, often known as a “buffer”. If the copy goes beyond the end of the buffer (“overruns” it), the application may unintentionally modify adjacent memory and potentially overwrite executable code or data, thereby causing a crash—or more sinisterly, overwrite its own code it with new executable actions to direct the computer to perform the bidding of the attacker.

This attack can be prevented if the incoming data is examined to ensure that it does not exceed a given size. However, failing to examine incoming data is typically not examined in this manner is a frequent programming mistake.

Malicious Content

Typically, Web pages contain both text and HTML (HyperText Markup Language) content that is generated by a server and interpreted by a client browser. Servers that generate static pages have full control over how the client will interpret the pages sent by that server. However, servers that generate dynamic pages do not have complete control over how their output is interpreted

1 by the client. If malicious content is introduced into a dynamic page, neither the
2 server nor the client has enough information to recognize that this has happened
3 and therefore take protective actions.

4 In HTML, to distinguish text from markup, some characters are treated in a
5 special manner. The grammar of HTML determines the significance of “special”
6 characters – different characters are special at different points in the document.
7 For example, the less-than sign (“<”) typically indicates the beginning of an
8 HTML tag. Tags can either affect the formatting of the page or introduce a script
9 program that the browser executes (e.g., the <SCRIPT> tag introduces code from a
10 variety of scripting languages).

11 Many Web servers generate Web pages dynamically. For example, a search
12 engine may perform a database search and then construct a Web page that contains
13 the results of the search. Any server that creates Web pages by inserting dynamic
14 data into a template should check to make sure that the to-be-inserted data does
15 not contain any special characters (e.g., “<”). If the inserted data contains special
16 characters, the user's Web browser is likely to mistake them for HTML markup.
17 Because HTML markup can introduce programs, the browser could interpret some
18 data values as HTML tags or script rather than displaying them as text.

19 The risk of a Web server not doing a check for special characters in
20 dynamically generated Web pages is that in some cases an attacker can choose the
21 data that the Web server inserts into the generated page. Then the attacker can
22 trick the user's browser into running a program of the attacker's choice. This
23 program will execute in the browser's security context for communicating with the
24 legitimate Web server, not the browser's security context for communicating with
25

1 the attacker. Thus, the program will execute in an inappropriate security context
2 with similarly inappropriate privileges.

3 This vulnerability is sometimes described as “cross-site scripting.”

4 Impact. Users may unintentionally execute scripts written by an attacker
5 when they follow untrusted links in Web pages, mail messages, or newsgroup
6 postings. Users may also unknowingly execute malicious scripts when viewing
7 dynamically generated pages based on content provided by other users.

8 Because the malicious scripts are executed in a context that appears to have
9 originated from the targeted site, the attacker may have full access to the document
10 retrieved (depending on the technology chosen by the attacker), and may send data
11 contained in the page back to their site. For example, a malicious script can read
12 fields in a form provided by the real server, and then send this data to the attacker.
13 The attacker can access restricted web sites from the client, or violate domain
14 security policies in other ways.

15 The specific impact can vary greatly depending on the language selected by
16 the attacker and the configuration of any authentic pages involved in the attack.
17 Some not immediately obvious examples are included here.

18 Attacks May Come in Many Forms. Attacks may be made persistent
19 through poisoned cookies. An attacker may access restricted web sites from the
20 client. Domain-based security policies may be violated. The use of less-common
21 character sets may present additional risk.

22 Input-Validation Filters

23
24 Since *unvalidated* input can cause these problems, then validating the input
25 prevents them from occurring. A set of instructions (either to a human or to a

computer) that examines input and rejects any invalid input is known as an input-validation filter. Input-validation filters for filtering input from computing components will prevent a malicious attacker from wreaking havoc in the manner described above. To validate an input, it may be filtered by the very application program that uses the input or it may be filtered before that program receives the input.

Fig. 1 shows a computer system, specifically a Web server 130, operatively coupled to a typical computer client 110 via a network 120, such as the Internet. The interloper on computer client 110 sends malicious input to the Web server 130 with hopes of infiltrating one of its applications (such as applications 134a, 134b, 134n) and/or its data.

With internalized filtering, the application programs themselves filter the input. With externalized filtering, the Web server 130 filters the input for the applications. More specifically, the input filter 132 of the server filters its input.

Internalized Filtering

Generally, software developers appear to be well positioned to write their code so that it filters incoming data to ensure that such data is valid and legal. When the application program filters its own input, then filtering is internalized. Unfortunately, most developers are focused on producing functional code rather than code resistant to such attacks; as a result, such security features are often ignored. In reality, software developers tend to be inconsistent in performing all necessary checks for validity of the inputs their applications receive.

Moreover, software developers may not actually be well positioned to write their code so that it filters incoming data to ensure that such data is valid and

1 legal. It is unrealistic to expect the developers to know every possible form of
2 attack on their software; new attacks are often invented, which lead to new
3 requirements for input validation. For example, the “format string attack,” as
4 described in <http://news.cnet.com/news/0-1003-200-2719802.html>, was
5 discovered in the summer of 2000. One approach to preventing this attack via
6 input validation requires rejecting all input involving percent signs. Therefore, it
7 is prudent practice to have a mechanism for performing additional validation
8 checks in addition to the internal checks.

9 10 Externalized Filtering

11 External filtering for an application is when input into an application is
12 filtered before the application receives the input. The filtering is performed
13 external to the application for which the input is destined.

14 The prudent practice of externalized filtering may be implemented by a
15 system administrator (or others). An administrator is well positioned to decide
16 when performing such input validation is worthwhile. A good rule of thumb is that
17 filtering should be done when performance considerations (e.g., the overhead
18 incurred by the validity checks) are outweighed by the explicit need for robustness
19 in environments, which cannot be guaranteed to be secure.

20 In general, externalizing the checks on the input data into a separate filter
21 (i.e., external to the application) ensures that no assumptions (about the
22 application’s input data) specified in the filter are left unwarranted.

23 As illustrated in dashed box 150 (labeled “Background”), many system
24 administrators, on their own, manually write one or more sets of filtering
25 instructions as needed for each application. A filter (such as input filter 132) uses

these instructions to filter incoming data. Sometimes the sets of filtering instructions are called “filter scripts.”

The system administrator (of box 150) typically writes these sets of filter instructions in one of several specific languages. Such languages may include those traditionally known as “scripting” languages. This human (such as a system administrator) may use a traditional text-based user-interface (UI), as shown by monitor 152. Effectively, the manually written instruction set 154 becomes the input filter 132 of the server 130. Typically, a server has a mechanism (e.g., ISAPI) for invoking external filters. This mechanism may be used to invoke the instruction sets as an input filter.

External Filters

In a typical computer system, external filters, such as filter 132, are little program modules are performed on the computer system when it is started. They stay in memory until the computer system shuts down. The external filters can be configured to receive a number of special filter-event notifications that occur with each request that the computer system receives, and with each response that the computer system generates in return. They are called external filters because they are external to the applications that actually receive the input. Common examples of such external filters are “ISAPI filters.”

ISAPI Filters. Internet Server API, an API for the Microsoft® Internet Information Server® (IIS) Web server. ISAPI filters are DLLs loaded into the process and they stay in memory until shut-down. After they are loaded, ISAPI filters can be configured to receive a number of special filter-event notifications

1 that occur with each input received (e.g., HTTP request that the Web server) and
2 with each response that the Web server generates in return.

3 When an ISAPI filter is loaded, the filter passes a data structure to the Web
4 server containing, in addition to other information, a bit field that specifies the
5 types of filter-event notifications for which the filter should be notified. Each time
6 one of these events occurs, an event notification is started, and every ISAPI filter
7 that is set to monitor that event is notified.

8 Instructions for Filtering

9
10 Herein, “instructions” are a set of commands that can be executed without
11 user interaction. Generally, they are computer-implemented instructions. A filter
12 language is a programming language through which one can write the instructions
13 of a filter.

14 To effect input-validation filtering, many humans (such as system
15 administrators) manually write filtering instructions, such as instruction set 154,
16 on an ad hoc basis. These instructions examine the incoming data as it is received
17 by the system, but *before* the destination applications receive that data. If the
18 incoming data passes muster, it is forwarded on to the applications that need it.

19 The following are examples of instruction sets (or partial sets) that a human
20 (such as a system administrator) may write to filter inputs:
21
22
23
24
25

C++ Example

```
BYTE IsBadChar[] = {
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0xFF,0xFF,0x00,0x00,0xFF,0xFF,0xFF,0xFF,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0xFF,0xFF,0x00,0xFF,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,0x00,
};

DWORD FilterBuffer(BYTE * pString, DWORD cChLen) {
    BYTE * pBad = pString;
    BYTE * pGood = pString;
    DWORD i=0;
    if (!pString) return (0);
    for (i=0; pBad[i]; i++) {
        if (!IsBadChar[pBad[i]])
            *pGood++ = pBad[i];
    }
    return (pGood - pString);
}
```

JavaScript Example

```
function RemoveBad(InStr){
    InStr = InStr.replace(/</g, "");
    InStr = InStr.replace(/>/g, "");
    InStr = InStr.replace(/"/g, "");
    InStr = InStr.replace(/'/g, "");
    InStr = InStr.replace(/%/g, "");
    InStr = InStr.replace(/;/g, "");
    InStr = InStr.replace(/(/g, "");
    InStr = InStr.replace(/)/g, "");
}
```

```

1      InStr = InStr.replace(/\&/g, "");
2      InStr = InStr.replace(/\+/g, "");
3      return InStr;
4  }

```

Perl Example

```

5      #! The first function takes the negative approach.
6      #! Use a list of bad characters to filter the data
7      sub FilterNeg {
8          local( $fd ) = @_;
9          $fd =~ s/[<>\"'\"%&();\|\\(\&\+)]//g;
10         return( $fd );
11     }
12
13     #! The second function takes the positive approach.
14     #! Use a list of good characters to filter the data
15     sub FilterPos {
16         local( $fd ) = @_;
17         $fd =~ tr/A-Za-z0-9\ //dc;
18         return( $fd );
19     }
20
21     $Data = "This is a test string<script>";
22     $Data = &FilterNeg( $Data );
23     print "$Data\n";
24
25     $Data = "This is a test string<script>";
26     $Data = &FilterPos( $Data );
27     print "$Data\n";

```

Shortcomings of Manually Generated Filter Instructions

Although the manual generation of input-validation sets of filter instructions does address the ills inflicted by the malicious attacks described above, it does so at the expense of additional manpower. Most humans (such as system administrators) are hard-pressed to find resources (e.g. time) to research and write customized instructions for each separate application receiving input and for each corresponding variety of malicious attacks.

This conventional, manual, ad hoc solution to the problem is slow and inefficient: each human (such as a system administrator) must manually write a

1 filtering instruction for each need. Such a programming practice is inherently
2 slow, tedious, and error-prone.

3 What is needed is an automated approach for generating filtering
4 instruction sets based upon what the human (such as a system administrator)
5 defines as valid input data and valid parameters of such data for a specific
6 application. This way, the human (such as a system administrator) can avoid the
7 time and expense of manually writing filtering code like the one shown above.

8 SUMMARY

9
10 Described herein is a technology for facilitating the automated generation
11 of input-validation software filters. At least one of the implementations, described
12 herein, provides a convenient graphical user interface (GUI). With this GUI, a
13 user (such as a human (such as a system administrator)) is able to quickly enter a
14 set of parameters defining valid inputs. Conversely, the parameters may define
15 invalid input. From the entered parameters, the implementation *automatically*
16 generates input-validation filters for filtering input from computing components.
17 With this implementation, the user does not manually generate filtering
18 instructions per se – she only specifies a high-level description of what should be
19 filtered, not how; thus, the user does not need to be familiar with any specific
20 filtering language.

21 This summary itself is not intended to limit the scope of this patent.
22 Moreover, the title of this patent is not intended to limit the scope of this patent.
23 For a better understanding of the present invention, please see the following
24 detailed description and appending claims, taken in conjunction with the
25

1 accompanying drawings. The scope of the present invention is pointed out in the
2 appending claims.

3 **BRIEF DESCRIPTION OF THE DRAWINGS**

4
5 The same numbers are used throughout the drawings to reference like
6 elements and features.

7 **Fig. 1** is a block diagram showing an example environment in which input
8 filters are employed. It shows a pictorial representation of the conventional
9 mechanism for manually creating filtering instructions. It also shows a pictorial
10 representation of an embodiment of the invention herein.

11 **Fig. 2** is a flow diagram showing an illustrative methodological
12 implementation of the invention herein.

13 **Fig. 3** is an example of a computing operating environment capable of
14 implementing an embodiment (wholly or partially) of the invention herein.

15 **DETAILED DESCRIPTION**

16
17 In the following description, for purposes of the explanation, specific
18 numbers, materials and configurations are set forth in order to provide a thorough
19 understanding of the present invention. However, it will be apparent to one skilled
20 in the art that the present invention may be practiced without the specific
21 exemplary details. In other instances, well-known features are omitted or
22 simplified to clarify the description of the exemplary implementations of present
23 invention, and thereby better explain the present invention. Furthermore, for ease
24 of understanding, certain method steps are delineated as separate steps; however,
25

these separately delineated steps should not be construed as necessarily order-dependent in their performance.

The following description sets forth one or more exemplary implementations of an Automated Generator of Input-Validation Filters that incorporate elements recited in the appended claims. These implementations are described with specificity in order to meet statutory written description, enablement, and best-mode requirements. However, the description itself is not intended to limit the scope of this patent.

The inventors do not intend these exemplary implementations to limit the scope of the present invention. Rather, inventors have contemplated that the present invention might also be embodied and implemented in other ways, in conjunction with other present or future technologies.

An example of an embodiment of an Automated Generator of Input-Validation Filters may be referred to as an “exemplary filter generator.”

Overview

The one or more exemplary implementations, described herein, of the present invention may be implemented (in whole or in part) by a filter generation system 170 and/or by a computing environment like that shown in Fig. 3.

The exemplary filter generator automatically generates input-validation filters for filtering input from computing components. In at least one implementation, it generates an input-validation filter from a set of automatically generated filtering instructions. These instructions are automatically generated from information provided by a user (such as a human (such as a system administrator)). That information identifies what constitutes the definition of valid

1 input data and valid parameters of such data for specific applications. Typically,
2 the user provides this information via a graphical user interface (GUI). What
3 constitutes valid input data and its valid parameters may be determined by the
4 user, provided by a specific application, and/or provided by a third party.

5 With the exemplary filter generator, filtering instructions are automatically
6 generated based upon a given set of assumptions on the parameters extracted from
7 the input. The parameter boundaries define how the incoming data is to be parsed.

8 A graphical UI defines both the parameter boundaries and the assumptions
9 each extracted parameter should satisfy. Using this graphical UI, the user defines
10 the parameter boundaries and assumptions for a to-be-generated filter instruction
11 sets.

12 These parameter boundaries and assumptions are persisted into description-
13 representation data structures. An XML data structure is an example of one such
14 description-representation data structure.

15 The filtering instructions are automatically generated from the persisted
16 description-representation data structures. Typically, the resulting filter
17 instructions are employed in the same manner as the conventional, manually
18 written filter instructions described above in the "Background" section. That is, the
19 data incoming into a specific application is filtered (using the filter instructions)
20 before the data is passed along to the application.

21 Two points at which the exemplary filter generator is particularly helpful is
22 1) during initial installation and setup of an application; and/or 2) when it is
23 necessary to quickly deploy a fix (i.e., a patch) to a security hole for an
24 application.
25

Applications for the resulting filters — which are automatically generated by the exemplary filter generator — include (but are not limited to) input-validation filtering, firewall filtering, and API (application programming interface) wrapper filtering.

Exemplary Filter Generating System

Fig. 1 shows the Web server 130 operatively coupled to the client 110 via the Internet 120. The interloper on client 110 sends malicious input to the Web server 130 with hopes of infiltrating one of its applications (such as applications 134a, 134b, 134n). With externalized filtering, such as that provided by input filter 132, the Web server 130 filters the input for the applications. The input filter 132 executes the set of filter instructions generated by the exemplary filter generator.

Fig. 1 shows the filter generation system 170. As its name suggests and like the conventional solution (shown in **Fig. 1** at background 150), the filter generation system 170 generates filter instructions, such as instruction 174, and loads those instructions into an externalized filter, such as input filter 132.

The filter generation system 170 includes four main components: a graphical user-interface (UI) 172; a description-representation (“DR”) synthesizer 176; a description-representation (“DR”) parser 182; and an instruction synthesizer 184. These components are listed following the data flow. User enters data at the graphic UI 172. Data flows through and is processed by the DR synthesizer 176, the DR parser 182, and the instruction synthesizer 184, in that order. Processed data, in the form of filter instructions, flows out of the instruction synthesizer 184.

For additional implementation details on the components of the exemplary filter generator and the description-representations, see the section below titled “Other Implementation Details”.

Graphical UI

The graphical UI 172 is a tool that aids in the generation of a description representation of both parameter boundaries and assumptions on parameters. With this UI 172, the user provides descriptions of various conditions that the inputs to a specific application need to satisfy. These descriptions may also be called “input-description-data.”

The descriptions define the properties of valid input and necessarily the converse - the properties (i.e., characteristics) of invalid input. More specifically, the descriptions define the parameter boundaries of incoming data to a given application and define the assumptions for those parameters (delimited by the boundaries). Although the defined descriptions could apply to multiple applications, they are primarily intended for filtering input to a specific application.

The user may obtain these descriptions from information provided with an application, provided by a third party (e.g., an independent software vendor, a trade magazine), and/or manually prepared by a user. Alternatively, this description may be automatically generated by examining available source code, if any. However, this option is typically not available to users because the source code of applications is typically not available.

In the exemplary filter generator, the graphical UI 172 is, generally speaking, a traditional easy-to-use graphic UI of the point-and-click variety.

1 Many operating systems provide an environment in which such a graphic UI may
2 be implemented. Examples of such operating systems include Microsoft®
3 Windows® XP, Microsoft® Windows® 2000, Microsoft® Windows® 98,
4 Microsoft® Windows® 95, Microsoft® Windows® ME, and Microsoft®
5 Windows NT® 4.0.

6 The graphical UI 172 of **Fig. 1** shows a representation of a window
7 containing traditional means of gathering information in such a UI. For example,
8 there may be radio buttons, drop-down lists, check boxes, data entry boxes, and
9 the like. Note that this technique is completely unlike the conventional instruction
10 writing technique illustrated in dashed box 150 of **Fig. 1**. It is more efficient,
11 quicker, easier, and less prone to errors than the conventional technique.

12 Description Representation Synthesizer

13
14 The DR synthesizer 176 takes the descriptions provided by the user via the
15 graphical UI 172 and generates “description representations,” which are
16 intermediate representations of the description provided by the user. This process
17 may also be called a transformation. The description representations represent the
18 descriptions of the parameter boundaries and assumptions provided by the user via
19 the graphical UI 172.

20 This may also be called “formal description.” In the exemplary filter
21 generator, the format for the description representations is in XML (eXtensible
22 Markup Language). More specifically, a small subset of XML may be employed.
23 This small subset may be particularly aimed at better readability (as text) even
24 without the use of any specialized tools or preliminary XML knowledge on the
25 user’s part.

1 The DR synthesizer 176 generates the description representations and
2 provides them to the filter instruction auto-generator 180. The exemplary filter
3 generator may temporally or permanently store the description representations in a
4 storage system. These stored description representations may be modified
5 manually as the user sees fit.

6 Alternatively, a user may manually create the description representations
7 independent of the graphical UI 172 and the DR synthesizer 176. These
8 independently created description representations may be sent to filter instruction
9 auto-generator 180.

10 Filter Instruction Auto-Generator

11
12 Collectively, the DR parser 182 and the instruction synthesizer 184 form the
13 filter instruction auto-generator 180. The auto-generator 180 outputs an
14 instruction, such as instruction 174. That instruction implements the filtering as
15 defined by the original descriptions provided by the user via the UI 172.

16 Collectively, the DR parser 182 and the instruction synthesizer 184
17 translate the description representations into a set of instructions. The details of the
18 implementations of these components depend heavily on the format of the
19 description representations and the chosen filtering language. With the
20 descriptions and examples provided herein, those of ordinary skill in the art can
21 implement these components. More specifically, it will take, at least, an order of
22 magnitude less time and effort for those of ordinary skill in the art to come up with
23 and implement these components, as compared to the situation where the
24 instruction has to be written manually in some special language.

1 The filter instruction auto-generator 180 may be pre-configured for each
2 application and each filter type. Consequently, the filter instruction auto-generator
3 180 interprets and translates the description representation into a given set of
4 filtering instructions on a given incoming parameter list. Therefore, the running
5 time of the exemplary filter generator does not add any overhead to the
6 application's execution time.

7 Alternatively, the exemplary filter generator may be implemented without
8 the DR synthesizer 176 and the DR parser 182. Instead, the filter instruction auto-
9 generator 180 may generate the instructions directly from the information
10 provided by the user via the graphical UI 172.

11 Automatically Generated Filter Instructions

12
13 The filter instruction auto-generator 180 generates the automatically
14 generated ("autogen") filter instructions 174. As shown in **Fig. 1**, the autogen filter
15 instructions 174 are sent to and effectively become the input filter 132. Upon the
16 receipt of input, the autogen filter instructions are executed *before* the actual
17 application receives that input. This performs the task of checking if all specified
18 assumptions (extracted from the given formal description) have been met and
19 takes appropriate steps (e.g., alerting the administrator, filtering out inappropriate
20 input).

21 Depending on the particular type of filter necessary, different filtering
22 languages may be applicable. For example, a filtering language often used in
23 server application space is VBScript, JavaScript, C, and C++. But, for other
24 configurations, other specially tailored (e.g., firewall, API wrapper) languages
25 may be utilized. For additional implementation details on the components of the

exemplary filter generator and the set of filter instructions, see the section below titled “Other Implementation Details”.

Input Filter

The filter instruction set delimits the parameters and defines the set of assumptions to test. When running the instruction set, the input filter decides whether the input satisfies the set of assumptions.

For example, assume that it is desirable to filter incoming URL requests. The filtering instruction set is registered as a COM object, and called by a “wrapper” ISAPI DLL when the URL is received. The instruction obtains the URL string. The (URL) string is then parsed into an array of parameters, according to the parameter boundary descriptions, as specified in the description representations and implemented in the set of instructions. Each of the parameters is tested to make sure it satisfies the assumptions, outlined in the description representations and implemented in the instructions. The output of the instructions specifies which assumptions (if any) have been violated.

If any assumptions have been violated, the filter may perform some kind of “filtering” action. For example, it may choose to refuse to pass the input to the application, or may modify the input to ensure that the assumption holds. It may send a notification to the system administrator. Alternatively, it may do anything else of the like.

Methodological Implementation of the Exemplary Filter Generator

Fig. 2 shows methodological implementation of the exemplary filter generator performed by the filter generation system 170 (or some portion thereof).

1 This methodological implementation may be performed in software, hardware, or
2 a combination thereof.

3 At 210, a user enters data thus defining the parameter boundaries and
4 assumptions. This data may be the “descriptions.” The user does this via a
5 graphical UI, such as the UI 172 of **Fig. 1**. At 212, the exemplary filter generator
6 synthesizes the descriptions to generate the description representations and these
7 are persisted (e.g., stored) at 214. Again, the exemplary filter generator employs
8 an XML data structure for the description representations.

9 At 216 of **Fig. 2**, the exemplary filter generator translates the description
10 representations into filter instructions. This translation typically includes a parsing
11 of the description representation and a synthesis of the filter instruction. This
12 process may be customized and streamlined by pre-configuring the translator for
13 known applications and filter types; for example, if the filter instructions are
14 described in a compiled language such as C++, this process may include
15 automatically invoking the compiler and linker on the filter instructions.

16 At 218, the just-generated filter instructions are loaded into and effectively
17 become the input filter, such as filter 132 of **Fig. 1**. In other words, the host system
18 using the filter feeds the input (acquired in block 210 of **Fig. 2**) into the filtering
19 instructions. Consequently, the instructions delimit the parameters and define the
20 set of assumptions to test. When running the instructions, the input filter decides
21 whether the input satisfies the set of assumptions. The process ends at 220.

22 **Exemplary Computing System and Environment**

24 **Fig. 3** illustrates an example of a suitable computing environment 900
25 within which an exemplary filter generator, as described herein, may be

1 implemented (either fully or partially). The computing environment 900 may be
2 utilized in the computer and network architectures described herein.

3 The exemplary computing environment 900 is only one example of a
4 computing environment and is not intended to suggest any limitation as to the
5 scope of use or functionality of the computer and network architectures. Neither
6 should the computing environment 900 be interpreted as having any dependency
7 or requirement relating to any one or combination of components illustrated in the
8 exemplary computing environment 900.

9 The exemplary filter generator may be implemented with numerous other
10 general-purpose or special-purpose computing system environments or
11 configurations. Examples of well known computing systems, environments, and/or
12 configurations that may be suitable for use include, but are not limited to, personal
13 computers, server computers, thin clients, thick clients, hand-held or laptop
14 devices, multiprocessor systems, microprocessor-based systems, set top boxes,
15 programmable consumer electronics, network PCs, minicomputers, mainframe
16 computers, distributed computing environments that include any of the above
17 systems or devices, and the like.

18 The exemplary filter generator may be described in the general context of
19 computer-executable instructions, such as program modules, being executed by a
20 computer. Generally, program modules include routines, programs, objects,
21 components, data structures, etc. that perform particular tasks or implement
22 particular abstract data types. The exemplary filter generator may also be practiced
23 in distributed computing environments where tasks are performed by remote
24 processing devices that are linked through a communications network. In a
25

distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

The computing environment 900 includes a general-purpose computing device in the form of a computer 902. The components of computer 902 can include, by are not limited to, one or more processors or processing units 904, a system memory 906, and a system bus 908 that couples various system components including the processor 904 to the system memory 906.

The system bus 908 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, such architectures can include an Industry Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA) local bus, and a Peripheral Component Interconnects (PCI) bus also known as a Mezzanine bus.

Computer 902 typically includes a variety of computer readable media. Such media can be any available media that is accessible by computer 902 and includes both volatile and non-volatile media, removable and non-removable media.

The system memory 906 includes computer readable media in the form of volatile memory, such as random access memory (RAM) 910, and/or non-volatile memory, such as read only memory (ROM) 912. A basic input/output system (BIOS) 914, containing the basic routines that help to transfer information between elements within computer 902, such as during start-up, is stored in ROM

1 912. RAM 910 typically contains data and/or program modules that are
2 immediately accessible to and/or presently operated on by the processing unit 904.

3 Computer 902 may also include other removable/non-removable,
4 volatile/non-volatile computer storage media. By way of example, **Fig. 3**
5 illustrates a hard disk drive 916 for reading from and writing to a non-removable,
6 non-volatile magnetic media (not shown), a magnetic disk drive 918 for reading
7 from and writing to a removable, non-volatile magnetic disk 920 (e.g., a “floppy
8 disk”), and an optical disk drive 922 for reading from and/or writing to a
9 removable, non-volatile optical disk 924 such as a CD-ROM, DVD-ROM, or other
10 optical media. The hard disk drive 916, magnetic disk drive 918, and optical disk
11 drive 922 are each connected to the system bus 908 by one or more data media
12 interfaces 926. Alternatively, the hard disk drive 916, magnetic disk drive 918, and
13 optical disk drive 922 can be connected to the system bus 908 by one or more
14 interfaces (not shown).

15 The disk drives and their associated computer-readable media provide non-
16 volatile storage of computer readable instructions, data structures, program
17 modules, and other data for computer 902. Although the example illustrates a hard
18 disk 916, a removable magnetic disk 920, and a removable optical disk 924, it is to
19 be appreciated that other types of computer readable media which can store data
20 that is accessible by a computer, such as magnetic cassettes or other magnetic
21 storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or
22 other optical storage, random access memories (RAM), read only memories
23 (ROM), electrically erasable programmable read-only memory (EEPROM), and
24 the like, can also be utilized to implement the exemplary computing system and
25 environment.

1 Any number of program modules can be stored on the hard disk 916,
2 magnetic disk 920, optical disk 924, ROM 912, and/or RAM 910, including by
3 way of example, an operating system 926, one or more application programs 928,
4 other program modules 930, and program data 932. Each of such operating system
5 926, one or more application programs 928, other program modules 930, and
6 program data 932 (or some combination thereof) may include an embodiment of
7 an input filter, an application program module, an input filter module, an
8 instruction, an interface, a filter-instruction automatic generator ("autogen"), a
9 synthesizer, and a memory.

10 A user can enter commands and information into computer 902 via input
11 devices such as a keyboard 934 and a pointing device 936 (e.g., a "mouse").
12 Other input devices 938 (not shown specifically) may include a microphone,
13 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and
14 other input devices are connected to the processing unit 904 via input/output
15 interfaces 940 that are coupled to the system bus 908, but may be connected by
16 other interface and bus structures, such as a parallel port, game port, or a universal
17 serial bus (USB).

18 A monitor 942 or other type of display device can also be connected to the
19 system bus 908 via an interface, such as a video adapter 944. In addition to the
20 monitor 942, other output peripheral devices can include components such as
21 speakers (not shown) and a printer 946 which can be connected to computer 902
22 via the input/output interfaces 940.

23 Computer 902 can operate in a networked environment using logical
24 connections to one or more remote computers, such as a remote computing device
25 948. By way of example, the remote computing device 948 can be a personal

1 computer, portable computer, a server, a router, a network computer, a peer device
2 or other common network node, and the like. The remote computing device 948 is
3 illustrated as a portable computer that can include many or all of the elements and
4 features described herein relative to computer 902.

5 Logical connections between computer 902 and the remote computer 948
6 are depicted as a local area network (LAN) 950 and a general wide area network
7 (WAN) 952. Such networking environments are commonplace in offices,
8 enterprise-wide computer networks, intranets, and the Internet.

9 When implemented in a LAN networking environment, the computer 902 is
10 connected to a local network 950 via a network interface or adapter 954. When
11 implemented in a WAN networking environment, the computer 902 typically
12 includes a modem 956 or other means for establishing communications over the
13 wide network 952. The modem 956, which can be internal or external to computer
14 902, can be connected to the system bus 908 via the input/output interfaces 940 or
15 other appropriate mechanisms. It is to be appreciated that the illustrated network
16 connections are exemplary and that other means of establishing communication
17 link(s) between the computers 902 and 948 can be employed.

18 In a networked environment, such as that illustrated with computing
19 environment 900, program modules depicted relative to the computer 902, or
20 portions thereof, may be stored in a remote memory storage device. By way of
21 example, remote application programs 958 reside on a memory device of remote
22 computer 948. For purposes of illustration, application programs and other
23 executable program components such as the operating system are illustrated herein
24 as discrete blocks, although it is recognized that such programs and components
25

1 reside at various times in different storage components of the computing device
2 902, and are executed by the data processor(s) of the computer.

3 4 **Computer-Executable Instructions**

5 An implementation of an exemplary filter generator may be described in the
6 general context of computer-executable instructions, such as program modules,
7 executed by one or more computers or other devices. Generally, program modules
8 include routines, programs, objects, components, data structures, etc. that perform
9 particular tasks or implement particular abstract data types. Typically, the
10 functionality of the program modules may be combined or distributed as desired in
11 various embodiments.

12 13 **Exemplary Operating Environment**

14 **Fig. 3** illustrates an example of a suitable operating environment 900 in
15 which an exemplary filter generator may be implemented. Specifically, the
16 exemplary filter generator(s) described herein may be implemented (wholly or in
17 part) by any program modules 928-930 and/or operating system 926 in **Fig. 3** or a
18 portion thereof.

19 The operating environment is only an example of a suitable operating
20 environment and is not intended to suggest any limitation as to the scope or use of
21 functionality of the exemplary filter generator(s) described herein. Other well
22 known computing systems, environments, and/or configurations that are suitable
23 for use include, but are not limited to, personal computers (PCs), server
24 computers, hand-held or laptop devices, multiprocessor systems, microprocessor-
25 based systems, programmable consumer electronics, wireless phones and

1 equipments, general- and special-purpose appliances, application-specific
2 integrated circuits (ASICs), network PCs, minicomputers, mainframe computers,
3 distributed computing environments that include any of the above systems or
4 devices, and the like.

5 6 **Computer Readable Media**

7 An implementation of an exemplary filter generator may be stored on or
8 transmitted across some form of computer readable media. Computer readable
9 media can be any available media that can be accessed by a computer. By way of
10 example, and not limitation, computer readable media may comprise “computer
11 storage media” and “communications media.”

12 “Computer storage media” include volatile and non-volatile, removable and
13 non-removable media implemented in any method or technology for storage of
14 information such as computer readable instructions, data structures, program
15 modules, or other data. Computer storage media includes, but is not limited to,
16 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
17 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
18 tape, magnetic disk storage or other magnetic storage devices, or any other
19 medium which can be used to store the desired information and which can be
20 accessed by a computer.

21 “Communication media” typically embodies computer readable
22 instructions, data structures, program modules, or other data in a modulated data
23 signal, such as carrier wave or other transport mechanism. Communication media
24 also includes any information delivery media.

1 The term “modulated data signal” means a signal that has one or more of its
2 characteristics set or changed in such a manner as to encode information in the
3 signal. By way of example, and not limitation, communication media includes
4 wired media such as a wired network or direct-wired connection, and wireless
5 media such as acoustic, RF, infrared, and other wireless media. Combinations of
6 any of the above are also included within the scope of computer readable media.

7 **Other Implementation Details**

8
9 This section describes additional details related to one or more exemplary
10 implementations of the present invention.

11 **Filtering Language**

12
13 VBScript (“Visual Basic Scripting Edition”) is the filtering language
14 employed by exemplary filter generator. VBScript is based on the Visual Basic®
15 (by the Microsoft Corporation) programming language, but is much simpler.

16 Of course, any other filtering language may be used, such as JavaScript.
17 For that matter, any form of language may be used that is acceptable to an input
18 filter, like filter 132 of Fig. 1.

19 **Representation Format**

20
21 XML (eXtensible Markup Language) is the description representation
22 format employed by the exemplary filter generator. XML is an emerging standard
23 for common representation of data in transit and it is especially touted for use in
24 distributed applications. It can be viewed as generic enough to allow the
25 representation of structures of any kind and for any purpose. It is extensible

1 instead of generating the appropriate set of filter instructions directly from the UI?
2 Compatibility and interoperability.

3 There may be a variety of ways in which the administrator may choose to
4 create or edit a filter description. Not strictly tying it to any particular UI, coupled
5 with the wide acceptance of XML as an emerging standard for data interchange,
6 are factors which allow for descriptions created on different systems running
7 different software to still be compatible. Administrators across platforms may
8 have access to different tools, yet they could choose their favorite way to
9 create/edit a filter description in XML and would not lose interoperability by
10 doing so. To avoid unnecessary changes to the UI even when new unwarranted
11 assumptions are unveiled, the UI is modularized and detached from the generation
12 of the XML description representation.

13 UI for Parsing Rules Descriptions. When administrators have the ability to
14 define the “meaning” of input parameters, they are able to exercise strict control
15 over which inputs are considered safe for an application. In the exemplary filter
16 generator, this is accomplished by observing the input not as a set of delimited
17 strings passed to an application, but rather as an unknown string of characters,
18 which is to be delimited first into what would be the parameters. Thus, part of the
19 skill goes into deciding how those parameters should be defined and deciding the
20 correct and full set of assumptions on all parameters.

21 With the exemplary filter generator, the system administrator can focus her
22 attention on deciding how those parameters should be defined rather than on how
23 to write instructions.
24
25

1 For example, consider the following input request, an incoming URL
2 request, to a web server:

3 `http://www.foomusic.com/my-query?page=%2Fpop%2F&club=classical`

4
5 Here is a possible break-up of this URL into parameters (for clarity we show them
6 separated by spaces):

7 `http // www.foomusic.com my-query page %2Fpop%2F club classical`

8 There are many ways in which an administrator can define (i.e., delimit) the
9 input parameters in a string like this. The more fine-grained the definitions are, the
10 larger the number and the smaller the size of the parameters. This in turn would
11 allow more precise handling of each separate parameter.

12 However, one can imagine defining other sets of parameters (presumably
13 more coarse-grained, aiming at capturing larger contexts) on the same input string.
14 In the end, when two or more (independent) sets of parameters are defined, along
15 with the assumptions on each one, the corresponding filtering instructions can be
16 executed on the given input string and the overall result would be the conjunction
17 of the individual instruction results. This technique can be used when the
18 complexity of the input string is expected to be high.

19 For example, here are two ways to define the boundary between one
20 parameter and the next: by specifying a single-character end-delimiter of the
21 current parameter, or by giving the exact length of that parameter. An exception to
22 this rule is the very last parameter – running to the end of the input string – it is
23 “delimited” by the NULL end-delimiter. (An end-delimiter character, if specified,
24 is excluded from the extracted parameter.)
25

1 UI for Description of Assumptions. The parameters need to satisfy a set of
2 assumptions. Desirable constraints may be expressed using a UI. In the description
3 of assumptions, each assumption is specified either on a parameter itself (*SELF*) or
4 on its size (*SIZE*). (The latter naturally only takes numerical values.)

5 The exemplary filter generator implements a set of thirteen relations (see
6 Table 1). More or less may be used by different implementations.

Relation name	Relational semantics
CONSISTS	The parameter value as a string consists of a specified set of characters, e.g. a-z, 0-9, etc.
EXCLUDES	The parameter value as a string excludes a specified substring
ENDS_IN	The parameter value as a string ends in a specified string (suffix)
LE	The parameter value / size is less than or equal to a specified numerical constant
LT	The parameter value / size is less than a specified numerical constant
GE	The parameter value / size is greater than or equal to a specified numerical constant
GT	The parameter value / size is greater than a specified numerical constant
EQ	The parameter value / size is equal to a specified numerical constant
LEX_LE	The parameter value as a string lexicographically precedes or is equal to a specified string
LEX_LT	The parameter value as a string lexicographically precedes a specified string
LEX_GE	The parameter value as a string lexicographically succeeds or is equal to a specified string
LEX_GT	The parameter value as a string lexicographically succeeds a specified string
LEX_EQ	The parameter value as a string is lexicographically equal to a specified string

Table 1. Example relations that may be used in defining assumptions on parameters

Herein, relations may have a disjunctive meaning (i.e., there can be more than one specified constant on the right-hand side, delimited appropriately) with the parameter required to satisfy the relation with at least one of the specified constants.

Consider, for example, the sample URL request and parameters provided above. Below is a sample list of assumptions (in natural language) that one might wish to impose on the first few parameters:

- Parameter 1 should consist of only lower-case alphabetical characters and should be “http” or “https”
- Parameter 2 should be exactly 2 characters long and these should be “//”
- Parameter 3 should be no more than 30 characters long, should consist of only lower-case alphabetical characters or the “period” character, and should have “.com” as suffix

XML Description Representations

The XML description of the assumptions consists of a set of clauses, each representing a single assumption. Semantically these clauses are AND-ed together to form a CNF (Conjunctive Normal Form) Boolean expression. Verifying that this expression is satisfied by the input data is equivalent to ensuring all stated assumptions have been met.

Representing assumptions naturally translates into CNF-expressions without the need for disjunction (OR) or negation (NOT) operations, which significantly simplifies the clause description in a formal language. There are a few notable exceptions, which do need a disjunction. For example, verifying if a parameter has its value—whether numeric or character string—among a list of given possibilities, and specifying a set of possible characters by which a parameter needs to be represented. All of these situations are handled without compromising the simplicity of the formal description format.

The formal description format described herein has been designed to be highly usable, striving toward a simpler and understandable XML file structure.

XML Parsing Rules Description. Below is a sample data structure of the XML description representations of parameter boundaries. XML parsing scheme description may be part of the results of using the exemplary filter generator to delimit the input request. It defines how input can be parsed into parameters using only a small uniform set of XML tags. Each parameter has its boundaries defined inside an instance of a `<param>` tag structure. This structure represents a collection of four substructures:

- `<number>` – the parameter number (in order from left to right);
- `<function>` – determining how the last (right-end) character of each parameter is determined – based on an exact parameter length in number of characters (`LENGTH`) or based on an end-delimiter character (`END_DELIM`);
- `<value>` – depending on the value of `<function>`, contains either the parameter length or the end-delimiting character;
- `<desc>` – (optional) concise natural-language description of the semantics of the current parameter.

Finally, the input to be parsed is merely a sequence of parameters, whose number has no a priori set upper bound. An example follows.

```
<input-parse>
  <param>
    <number>1</number>
    <function>END_DELIM</function>
    <value>:</value>
    <desc>Protocol name</desc>
  </param>
  <param>
    <number>2</number>
```

```

1      <function>LENGTH</function>
2      <value>2</value>
3      <desc>Double slash, i.e. //</desc>
4      </param>
5      <param>
6          <number>3</number>
7          <function>END_DELIM</function>
8          <value>/</value>
9          <desc>Server name</desc>
10     </param>
11     <!-- The definitions of parameters 4-7 have been omitted for brevity -->
12     <param>
13         <number>8</number>
14         <function>END_DELIM</function>
15         <value>NULL</value>
16         <desc>Second (real) parameter value</desc>
17     </param>
18 </input-parse>

```

Reading the above XML description, parameter 1 (semantically, the protocol name) starts from the beginning of the input string and ends before the first colon character. Then, parameter 2 (meant to be the delimiting double slash after the protocol name) picks up immediately after the colon and is exactly 2 characters long. Parameter 3 (supposed to contain a server name) starts immediately after parameter 2 and continues until the next slash-character in the input, etc. Parameter 8 (in the sample meant to hold the value of the “club” property) starts (naturally) after parameter 7 and is comprised of all characters until the end of the input string is reached.

Description of Assumptions. Below is a sample data structure of the XML description representations of assumptions — the output from the corresponding UI on the same sample URL request. Various assumptions on input data (parameters) could again be described using a very limited uniform set of XML tags.

```

22     <param>
23         <number>1</number>
24         <func>SELF</func>
25         <cond>CONSISTS</cond>
26         <delim>NULL</delim>
27         <value>a-z</value>
28         <desc>Parameter 1 should consist of only lower-case alphabetical
29         characters</desc>

```

```

1      </param>
2      <param>
3          <number>1</number>
4          <func>SELF</func>
5          <cond>LEX_EQ</cond>
6          <delim>,</delim>
7          <value>http,https</value>
8          <desc>Parameter 1 should be one of the strings "http" or "https"</desc>
9      </param>
10     <param>
11         <number>2</number>
12         <func>SIZE</func>
13         <cond>EQ</cond>
14         <delim>NULL</delim>
15         <value>2</value>
16         <desc>Parameter 2 should be exactly 2 characters long</desc>
17     </param>
18     <param>
19         <number>2</number>
20         <func>SELF</func>
21         <cond>LEX_EQ</cond>
22         <delim>NULL</delim>
23         <value>//</value>
24         <desc>Parameter 2 should be the string "//"</desc>
25     </param>
26     <param>
27         <number>3</number>
28         <func>SELF</func>
29         <cond>CONSISTS</cond>
30         <delim>,</delim>
31         <value>a-z,</value>
32         <desc>Parameter 3 should consist of lower-case alphabetical characters and
33         "."</desc>
34     </param>
35     <param>
36         <number>3</number>
37         <func>SIZE</func>
38         <cond>LE</cond>
39         <delim>NULL</delim>
40         <value>30</value>
41         <desc>Parameter 3 should be at most 30 characters long</desc>
42     </param>
43     <param>
44         <number>3</number>
45         <func>SELF</func>
46         <cond>ENDS_IN</cond>
47         <delim>NULL</delim>
48         <value>.com</value>
49         <desc>Parameter 3 should end with the string ".com"</desc>
50     </param>

```

As the above example illustrates, each assumption on a parameter is described within a separate `<param>` tag structure. This structure is comprised of six tags:

- `<number>` – the number (in order from left to right) of the parameter to which the assumption applies;

- `<func>` – the parameter function, determining whether the current assumption concerns the parameter value (`SELF`), or its size in number of characters (`SIZE`);
- `<cond>` – the relation between parameter value / size (depending on the function) and the list of constants specified in the `<value>` tag;
- `<delim>` – a single character delimiting constants in the `<value>` list (if the list contains a single constant, a `NULL` delimiter should be specified);
- `<value>` – a list of one or more constants being related to the parameter value / size by the function;
- `<desc>` – (optional) concise natural-language description of the semantics of the current assumption.

Rather than include a `<value-type>` tag in the XML structure, separate relations are defined for each type domain of possible arguments to simplify the design and avoid ambiguities. For example, `GT` expects a numerical argument (an integer or a floating-point number), while `LEX_GT` expects a character string to be lexicographically compared to the parameter.

If the `<value>` list contains more than one constant, all subsequent constants should be of the same type as the first one. For instance, the relation `EQ` expects one or more numerical values, all necessarily of the same type. An advantage of using a GUI to generate the intermediate XML descriptions is that this approach gives the benefit of automatic type checking at description-generation time.

When there are two or more constants in the list, the relation has the semantics of a disjunction between a set of relations, each one between the parameter value / size and the next constant in the list. In other words, the

parameter in question satisfies the specified relation with the list if it does so with at least one of the constants in the list. For instance, the second assumption (in the XML above) demands that parameter 1 be lexicographically equal to one of `http` or `https`, while the fifth assumption states that parameter 3 should consist of characters from the set of letters `a-z` and the `."` character.

These two options—checking if a parameter has its value among a list of values, and specifying the allowable character set for representing a parameter—are the only situations when a disjunction operator is involved in specifying assumptions. Even so, they can be easily handled in the implementation without compromising the simplicity of the model, in which each assumption corresponds to exactly one clause and the set of assumptions is a set of clauses joined by conjunction.

The constants in the `<value>` list are represented according to the rules of XML; e.g., an ampersand character should be `&`, a less-than character should be `<`, etc. The exemplary filter generator facilitating the description generation takes care of this.

Server Filter

Filters are described in a structure such as `<server-filter>`. It specifies:

- `<param-count>` – the total number of parameters;
- `<max-total-length>` – the maximum allowed length of the application input (as it is given initially before being parsed into its constituent parameters);
- `<param-desc>` – (optional) concise natural language description of the parameter semantics;

- one or more `<param>` structures (as described above), each one describing an assumption on some parameter;
- zero or more `<complex-cond>` structures, each one “pointing” to an XML file and an optional description (a more in-depth discussion follows shortly).

Here is an example of a `<server-filter>` structure:

```
<server-filter>
  <param-count>8</param-count>
  <max-total-length>100</max-total-length>
  <param-desc>See the XML parsing schema descriptions</param-desc>

  <!-- Definitions of assumptions on parameters are omitted here -->

  <complex-cond>
    <xml-file>MoreComplex.xml</xml-file>
    <cond-desc>Some more complex assumptions</cond-desc>
  </complex-cond>
</server-filter>
```

Complex assumptions are specified in dedicated XML files, referred to by an `<xml-file>` tag inside the `<complex-cond>` structure (as in the above example). They are a special provision for making the model extensible by adding a level of indirection, which saves us from having to change the core XML format and/or the filter-generating code. They also allow any additions to the filter formats to happen incrementally (i.e., as more reports and exploits become known) and with ensured backward compatibility. Thus, assumptions which otherwise would not fit into the above-described structure can still be described, albeit in separate XML files.

Filter Generator

The filter instruction generator (such as filter instruction auto-generator 180) parses the input (XML) descriptions of parameter boundaries and assumptions on those parameters, and (depending on the type of filter, e.g. <server-filter> denotes back-end server filters) synthesizes a set of instructions (in an appropriate language), which can then be executed by a filter on an incoming input string in order to determine if the assumptions on the specified set of parameters all hold.

The instruction set generation process of the exemplary filter generator is completely detached from the real-time processing of incoming requests of the filter. In other words, the instruction set can be generated entirely offline; and thus performance is not an issue for the generation.

Filtering Instructions

A filtering instruction set is registered as a COM object, then invoked by the server via an ISAPI DLL. Alternatively, the filtering instruction set may be the index.html file, which a browser opens first when a URL comes in.

The resulting filter obtains the input (URL) line and parses it into a set of parameters. Assumptions are then verified against the delimited parameters and a conclusion is reached on whether all of the assumptions have been satisfied or whether some have failed and which ones, if so.

Exemplary Applications of the Exemplary Filter Generator

One application of the exemplary filter generator is in providing express response to security breaches through the dissemination of filtering instruction sets while the next set of security patches are still under development. This can save time, money, and reputation for the vendor of the application under attack. In many cases, where a security breach involves a simple omission of an input verification—and these have accounted for a large fraction of recent attacks—it could literally take a security expert minutes to come up with the right set of parameters and assumptions on them to generate a filtering script.

Furthermore, this can be accomplished without knowing the details of the application's source code and even without having access to it. Consequently, not only the application vendor will be able to supply filters, although there is value in knowing that a specific filtering script was distributed by a trusted party. The resulting script can ensure that attempts to exploit the same vulnerability will be fended off reliably in the future.

As a comparison, it can take weeks and even months (of analyzing, developing and testing) before a reliable patch to a widely deployed commercial software product is ready for distribution. In the case of legacy software, no longer supported by its original vendor, patches will likely never come out; however, filters would be quite easy and inexpensive to create and distribute.

Another application is the automatic generation of firewall filters. Generalized firewall configuration languages have been proposed; a further step might be to automate the generation of configuration instructions in such a language using GUI-based tools.

1 Still another application is the encapsulation of existing API function
2 libraries. The idea is that the wrapper filter would intercept the calls to its library
3 routines, do the necessary validation checks and only "forward" the calls to the
4 corresponding functions if it is safe to do so. When a vulnerability is found, the
5 appropriate wrapper filter could be generated and applied as a stopgap until an OS
6 or application patch is fully tested and released. In fact, application programmers
7 could perhaps specify at development time the assumptions made by their code
8 about its inputs, and generate the appropriate filters themselves. Then an
9 administrator could decide, based on the hostility of the environment, the
10 sensitivity of the application, and the performance constraints on the system,
11 whether the filter should be installed to improve security or omitted to address
12 performance.

13 Conclusion

14
15 Although the invention has been described in a language specific to
16 structural features and/or methodological steps, it is to be understood that the
17 invention defined in the appended claims is not necessarily limited to the specific
18 features or steps described. Rather, the specific features and steps are disclosed as
19 preferred forms of implementing the invention.
20
21
22
23
24
25